



EE 122: Introduction To Communication Networks

Fall 2012

Scott Shenker

<http://inst.eecs.berkeley.edu/~ee122/>

Materials with thanks to Jennifer Rexford, Ion Stoica, Vern Paxson
and other colleagues at Princeton and UC Berkeley

Are you in the right 122 class?

- **Spring** offering: taught by **EE** faculty
 - More emphasis on diverse link technologies, wireless, communication theory, and mathematical analysis
- **Fall** offering: taught by **CS** faculty
 - More emphasis on Internet architecture and real-world practice
- Classes are very different in content and style

Is 122 the right class for you?

- Want to understand the “why” of networking?
 - Not just looking for definitions and techniques
- Ready for some fun?
 - Are you willing to laugh at my bad jokes
- Willing to actively participate in class?

What is “Active Participation”?

- Ask and answer questions
 - Not just the same ten students
- Participate in class “exercises”
 - We will act out routing, do joint design tasks, etc.
- Sit towards the front
 - Room is way too large
- Go without electronic access for almost 90 minutes
 - Put all laptops/phones/etc away, at least for today
 - You’ll have a 5 minute break in the middle to get online

Today's lecture will cover two topics

- Course overview
 - Material covered
 - People involved
 - Policies and administrivia

5 Minute Break

- Four basic questions about networking
 - Why are networking courses so terrible?
 - Why is it important to study networking?
 - Why is this an exciting time for networking?
 - Why is networking so hard?

You might not understand this lecture

- My jargon may be unfamiliar
 - Packets, hosts, etc.
- Don't worry, you'll pick it up soon enough
 - And you won't have missed anything in the mean time

What is a course on networking?

- There are many networks
 - Telephone (landline) networks
 - Cellular networks
 - Supervisory control and data acquisition networks
 - Frame relay networks
 - Optical networks
 -
- **We won't study any of them....**

Class will focus almost exclusively on the Internet

Networks versus “The Internet”

- The Internet is not a particular kind of network
 - It is not a battle between, say, Ethernet and Internet
- The Internet ties different networks together
 - The Internet
- ***Why does this matter?***

Goals for a network technology

- Speed
- Cost
- Port-density
- Reliability
- Other “features”
 - Quality of service, security, etc.
-

Goals for the Internet

- Ability to connect many different networks
- Ability to scale to entire world
- Ability to recover from failures
-

These are harder and more interesting goals!

(more architectural than engineering)

Architecture vs Engineering

- Architecture:
 - The allocation of functionality and definition of interfaces among elements
- The Internet “architecture” is the decision about what tasks get done, and where:
 - In the network, or in the hosts
 - Engineering is more about how tasks get done
- These architectural decisions play a crucial role in scaling, heterogeneity, robustness, etc...
 - **This is what I spend my life worrying about**

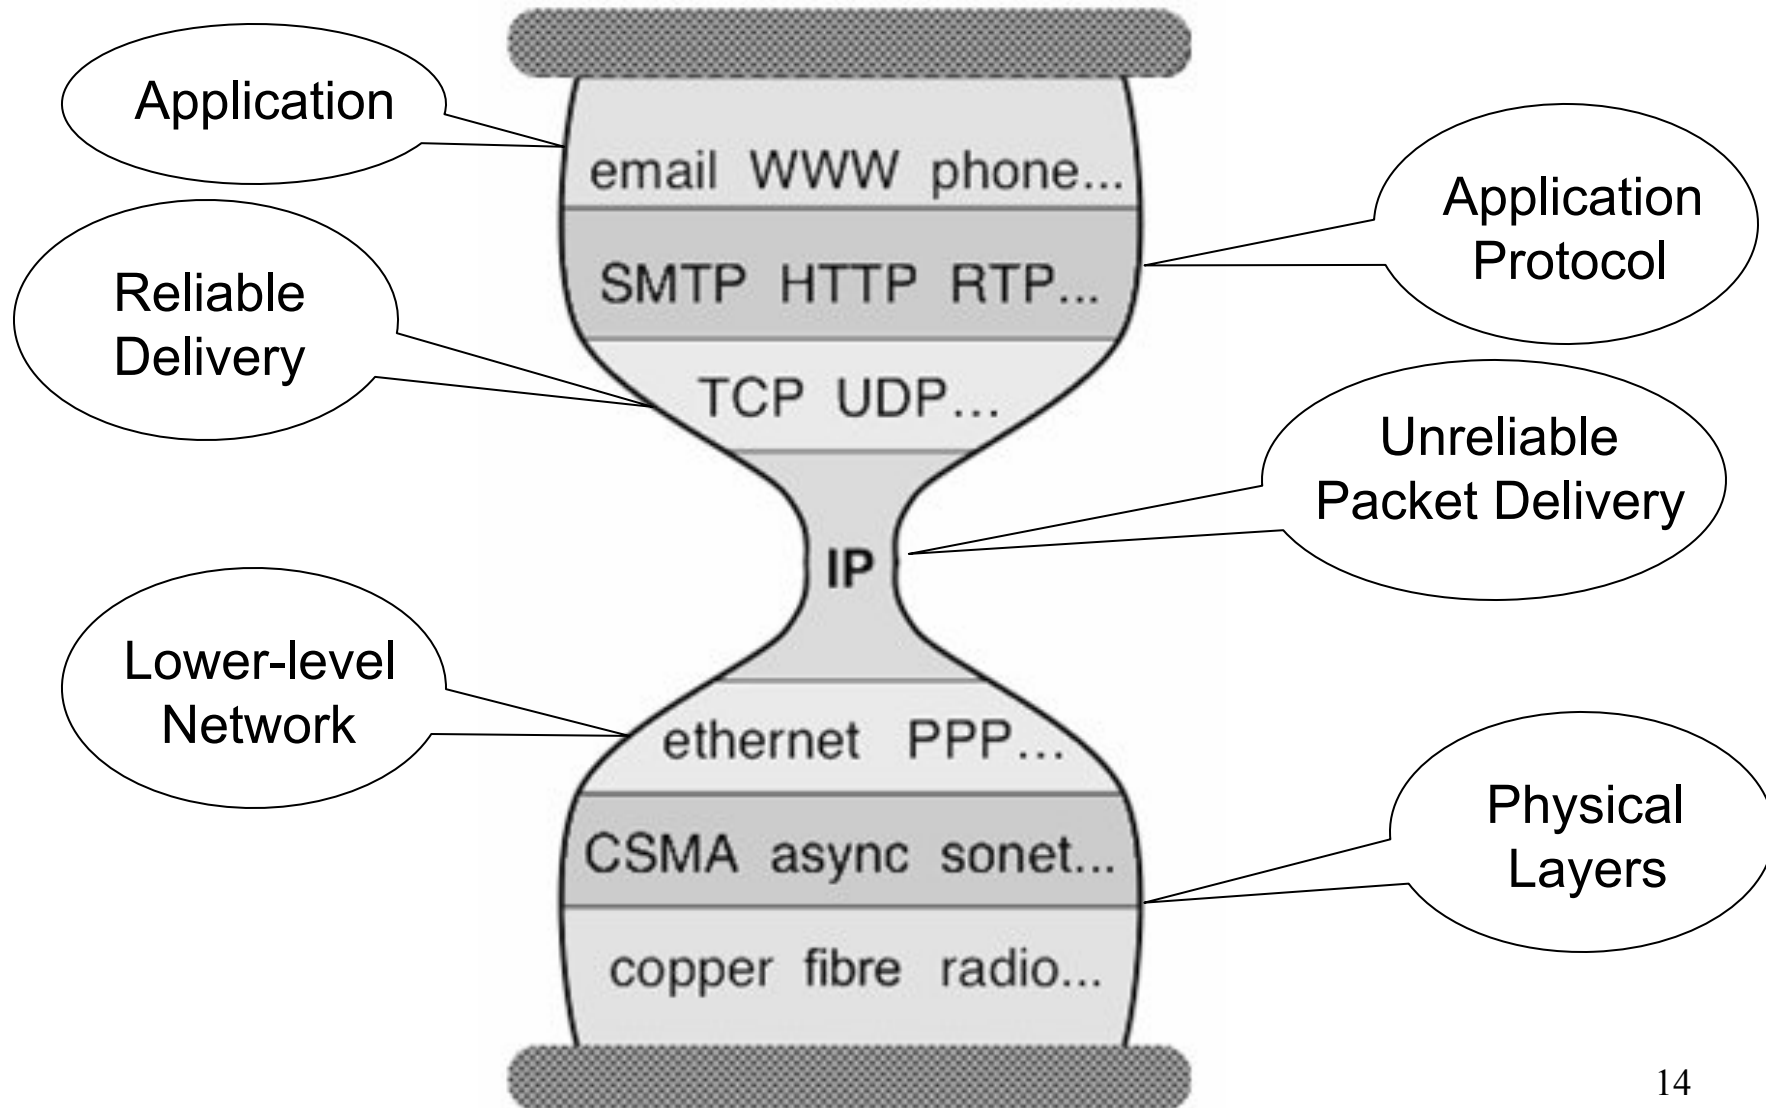
What topics will course cover?

- The core of the Internet “architecture”:
 - IP, DNS, BGP
- Other technologies crucial to the Internet
 - Higher-level protocols: TCP, HTTP....
 - Crucial lower-level technologies: Ethernet, wireless...
 - These are the two network technologies we will study because they raise interesting questions about shared media
- Won't cover network topics not crucial to Internet
 - But that doesn't mean they aren't interesting
 - E.g., sensornets, low-level encoding, radio technology

Various perspectives on Internet

- Different levels of abstraction
 - Basic concepts versus actual protocols
- Different geographic scales:
 - LAN vs Enterprise vs WAN vs Interdomain
- Different conceptual approaches:
 - Architecture vs Protocol vs Algorithm
- Different aspects of functionality:
 - Different “layers” focus on different tasks

The Internet: an hourglass *with layers*



Most networking courses

- Organized around layers:
 - Top-down (K&R) *[book we are using]*
 - Bottom-up (P&D)
- Why not for this course?
 - Main distinction is not where functionality is implemented
 - It is between **basic concepts** and **actual realization**
 - If you walk through layers sequentially, do both at once
- I care most about teaching the concepts
 - Implementations needed to put these ideas into practice
 - But don't want to lose basic concepts in sea of details


First half of course: Basics

- General overview (3 lectures after today)
 - Packet switching, basic design principles
- Idealized view of network (3 lectures)
 - Focus on fundamental conceptual questions
 - Ignore all real-world unpleasantness
- Making this vision real (5 lectures)
 - IP, TCP, DNS, Web
 - Emphasize concepts, but deal with unpleasant realities

Fundamental conceptual questions

- How can you deliver packets from source to destination?
- How do you build reliable transport on top of an unreliable network?
- How can you federate a set of competing ISPs?
-

Second half of course: Various topics

- Congestion control
 - Advanced topics in routing
 - Multicast and QoS
 - Security
 - Ethernet
 - Wireless
 - Software-defined networking
 - Alternate architectures
- Multiple Access**
- 

People: Teaching Assistants

- Anand Iyer
- Andrew Or
- Aurojit Panda
- Colin Scott
- Gautam Kumar
- Kay Ousterhout
- Thurston Dang
- Shaddi Hasan
- Tathagata Das

Instructor: Scott Shenker

- Trained as a physicist (phase transitions, chaos)
- Research: physics, economics, operating systems, security, distributed systems, datacenter design...
 - Diversity reflects my learning and teaching style
- For last 25 years, main focus has been networking and Internet architecture
- Office hours Thursday 2:00-3:00 in 415 Soda Hall
 - **Always** ping me by email before heading over
 - And by appointment (arrange by email)
 - On campus M, T, Th; on email 24 hours/day
 - **Available after class**

My teaching style is not for everyone...

- Next few slides provide a small taste of my flaws
- With a few comments from my 2010 class evals

I won't remember your name

- Prosopagnosia (as described by Oliver Sacks)
- In my case, it isn't recognizing faces, but attaching names to faces
- Don't take it personally....
 - Can't attach names to faces for over 50% of the faculty

I don't think visually

- “Uses blackboard terribly. Very poor diagrams when using it. and not legible also.”
- “For the love of god, use more pictures and diagrams.”
- I'm not going to turn into a blackboard virtuoso or animation wizard
- Ask TAs for pictures
- Will try to use other visual means
 - Watch for our re-enactment of routing.....

When you look bored, I speed up

- “Pace gets faster if no one asks questions.”
- If you are bored, feel free to sleep (*at your peril*)
- If you are lost, ask me a question!
 - Or just yell “HELP!”

I hate details

- “Moves very quickly during difficult topics and slowly during basic topics.”
- Will try to go over examples in more depth
- Sections will go over examples in even more depth

Can't always engage class

- “He asks questions but no one answers”
- Will try various approaches to get you to talk
- **But, I don't ask questions to get answers.....**

I ask questions so you can think!

- The pause after I ask a question is the only time you get to think
 - When I ask a question, I don't care if you answer it
 - But please, ***think about the question!***
- The best way to understand networking is to *first* try to solve the design issues yourself
 - Then the current solution will make a lot more sense
- **Internet not principled design, mostly ad hoc**
 - Can't “follow the logic”, have to try designing it yourself

Administrivia: Textbook

- J. Kurose and K. Ross, *Computer Networking: A Top-Down Approach*, 6th Edition, 2012.
 - 5th Edition ok, but translate the reading assignments
- For reasons I will discuss later, networking is a very hard area to teach. The textbook isn't great, but it is about as good as they come.
- Use only as reference, and source of examples
 - Those details I like to ignore? Go read about them.
- **You will not be tested on material I didn't cover** 28

Three projects

- Project 1: Reliable transport (in simple simulator)
- Project 2: Routing (in simple simulator)
- Project 3: Adding functionality to a home router
 - Larger project, in two phases
 - Will implement on your own Plug computer
 - Donated by Marvell
- **TAs will handle all project-related questions!**

Additional Lectures?

- Stanford is starting an online networking course
- I may assign some lectures as background

Class communications

- Web site: <http://inst.eecs.berkeley.edu/~ee122/>
 - Assignments, lecture slides
 - Please don't use slides to answer questions I ask
- Use bspace to hand in homework, send announcements
- Use Piazza for all other intraclass communication
 - You should all be signed up now
- Fill out questionnaire!
 - <http://tinyurl.com/8ererxf>

Did you get my email yesterday?

- If not, then either:
 - You aren't yet on our bspace class list, or
 - Your email address on that list is incorrect, or
 - There is some other failure mode (spam, etc.)
- Please send me email ASAP if you did not get that email from me.

Who Are You? (so far)

- 58% seniors, 34% juniors
- 30% love networking, 40% just looking for credits
- 13% no proficiency in python
- 75% have written programs > 1000loc
- 29% have taken 162, 12% never plan on taking it!
- 69% got the limit wrong, 74% got the coins right
- Varying levels of network familiarity
 - 60% know IP, 55% know DNS, 1% know BGP,....

Class workload

- Three projects (covered earlier)
- Four homeworks
 - Strict due dates (no slip days!)
 - Deadlines are generally 5:00PM prior to lecture
 - **Deadly boring, but designed to prepare you for exams**
 - **May also distribute optional worksheets (not graded)**
- Exams
 - Midterm: **Tuesday October 9** in class
 - Final: **Thursday Dec 13** location TBD, 11:30AM-2:30PM
 - **Closed book, open crib sheet**

Grading

Homeworks	20% (5% each)
Projects	40% (10+10+20)
Midterm exam	15%
Final exam	25%

- Course graded to mean of B

Participation Requirement

- Must speak up in class, or see me in office hours
 - At least once, or else you flunk. Period.
- If you've asked or answered a question, send email to your TA that day repeating your question or answer. Use emails of the form:
 - ee122.name@gmail.com
 - {anand, andrew, colin, gautam, kay, panda, thurston}
- If you've seen me in office hours, send email to me summarizing what we talked about (ee122.scott)

No Cheating

- Fine to *talk* with other students about assignments
 - But only general concepts, not specifics
- General rule: no copying of specifics
 - If you're unsure, then ask.
- Will use automated similarity detection
- *Don't be an idiot....*

5 Minute Break

Questions Before We Proceed?

And just a quick word before part 2...

- How many of you are eligible to vote?

Four Questions

- **Why are networking courses so terrible?**
- **Why is it important to study networking?**
- **Why is this an exciting time for networking?**
- **Why is networking so hard?**

1: Why are networking courses so bad?

- Reason 1: The basic Internet architecture has not changed since its invention over 35 years ago
 - Even IPv6 is very similar to IP
- Can't test an Internet architecture in lab or testbed
 - So we only understand what we currently have
- **We are teaching history, not principles**
 - You will learn “first tries” not “fundamental answers”
 - As if we taught MS-DOS in an operating system course

Bad networking courses, continued....

- Reason 2: No intellectual framework for networking
- Internet inventors defined a brilliant paradigm
 - Since then, community has focused on protocols to realize this paradigm
- Research community has failed to provide a general framework for understanding protocols
- We therefore just teach a big bag of protocols
 - And let you try to make sense of it yourself

Reason 3: Quote from John Day

There is a tendency in our field to believe that everything we currently use is a paragon of engineering, rather than a snapshot of our understanding at the time. We build great myths of spin about how what we have done is the only way to do it to the point that our universities now teach the flaws to students (and professors and textbook authors) who don't know better.

I will try to overcome these problems

- Focus when possible on “fundamental questions”
 - And will present alternative designs in a few lectures
- You will have to learn the current design
 - But I will point out where it falls short
- You will end up with a mixture of the “big picture” and “current design details”

2: Why important to study networking?

- Huge impact
- New paradigm
- Unresolved challenges

Internet has had tremendous impact

- Internet changed the way we gather information
 - Web, search engines
- Internet changed the way we relate to each other
 - Email, facebook, twitter
- **Which would you choose?**
 - Computers without the Internet (standalone PCs)
 - Internet without modern computers

The Internet introduced new paradigm

- Completely different from the phone network
- Inventors had to overcome strong technical and commercial resistance to realize their dreams
 - Motivation not for personal gain, but societal benefit!
- A true success story of “thinking differently”
 - Their strong vision kept the design on track
 - Brilliant in conception, sometimes weak in execution
- While mired in details, ***leave room for awe***

Many challenges remain unsolved

- Security
 - Security of infrastructure
 - Security of users
- Availability
 - Internet is very resilient
 - But availability is not sufficient for critical infrastructures
- Evolution
 - It is too hard to change the Internet architecture

3: Why an exciting time in networking?

- The “architecture” won’t change
 - But how we build and manage networks will
- Industry has been closed, stagnant, and feudal
- **But we are on the verge of a revolution!**
 - Commodity hardware making inroads
 - Developing intellectual (and practical) framework of applying systems principles of abstraction and modularity
- Full disclosure: I had a startup in this area
 - But approach endorsed by almost everyone else

4: Why is Networking Hard?

- There are many challenges that make designing the Internet harder than just passing bits on a wire
- *Which of these apply to the phone network?*

Scale

- Over 2 Billion Internet users.....

WORLD INTERNET USAGE AND POPULATION STATISTICS December 31, 2011

World Regions	Population (2011 Est.)	Internet Users Dec. 31, 2000	Internet Users Latest Data	Penetration (% Population)	Growth 2000-2011	Users % of Table
<u>Africa</u>	1,037,524,058	4,514,400	139,875,242	13.5 %	2,988.4 %	6.2 %
<u>Asia</u>	3,879,740,877	114,304,000	1,016,799,076	26.2 %	789.6 %	44.8 %
<u>Europe</u>	816,426,346	105,096,093	500,723,686	61.3 %	376.4 %	22.1 %
<u>Middle East</u>	216,258,843	3,284,800	77,020,995	35.6 %	2,244.8 %	3.4 %
<u>North America</u>	347,394,870	108,096,800	273,067,546	78.6 %	152.6 %	12.0 %
<u>Latin America / Carib.</u>	597,283,165	18,068,919	235,819,740	39.5 %	1,205.1 %	10.4 %
<u>Oceania / Australia</u>	35,426,995	7,620,480	23,927,457	67.5 %	214.0 %	1.1 %
<u>WORLD TOTAL</u>	6,930,055,154	360,985,492	2,267,233,742	32.7 %	528.1 %	100.0 %

Dynamic Range

- Round-trip times (**latency**) from 10 μ secs to secs
– 5 orders of magnitude
- Data rates (**bandwidth**) from kbps to 100 Gbps
– 8 orders of magnitude
- **Queuing** delays in the network vary from 0 to secs
- **Packet loss** varies from 0 to 90+%
-

Diversity of end systems

- *Cell phones*
- *Supercomputer clusters*
- *Tablets*
- *Televisions*
- *Gaming consoles*
- *Web cams*
- *Automobiles*
- *Sensing devices*
- *Picture frames*
- *Security systems*
- *Power grid*
- *.....*

Diversity of application requirements

- Size of transfers
- Bidirectionality (or not)
- Latency sensitive (or not)
- Tolerance of jitter (or not)
- Tolerance of packet drop (or not)
- Need for reliability (or not)
- Multipoint (or not)
-

Ad hoc deployment

- Can't assume carefully managed deployment
 - Network must work without planning

Networks contain many components

Links



Fibers



Coaxial Cable

Interfaces

Ethernet card



Wireless card



Switches/routers

Large router



Telephone switch



They can all fail....

- Consider communication that uses 50 components
 - Assume each work correctly 99% of the time
 - What is likelihood communication fails?
- Answer: success requires that they all function, so failure probability = $1 - (.99)^{50} \approx 39.5\%$
- Even if nodes are 99.9% reliable, failure probability is still close to 5%
- **Must design the system to expect failure!**
- Joke: Why is the Internet like a 12-step program?

Greed

- There are greedy people out there who want to:
 - Steal your financial information (bank, credit card, etc.)
 - Use your computer for attacks
- There is a thriving underground economy for compromised computers and financial information

6 controlled.

7 71. ANCHETA would develop a worm which would cause infected
8 computers, unbeknownst to the users of the infected computers, to:

9 a. report to the IRC channel he controlled;

10 b. scan for other computers vulnerable to similar
11 infection; and

12 c. succumb to future unauthorized accesses, including
13 for use as proxies for spamming.

20 his worm caused 1,000 to 10,000 new bots to join his botnet over
21 the course of only three days.

18 73. ANCHETA would then advertise the sale of bots for the
19 purpose of launching DDOS attacks or using the bots as proxies to
20 send spam.

21 74. ANCHETA would sell up to 10,000 bots or proxies at a
22 time.

23 75. ANCHETA would discuss with purchasers the nature and
24 extent of the DDOS or proxy spamming they were interested in

9 79. ANCHETA would accept payments through Paypal.

15 103. In or about August 2004, ANCHETA updated his
16 advertisement to increase the price of bots and proxies, to limit
17 the purchase of bots to 2,000 "due to massive orders," and to warn,

14 adware on those computers without notice to or consent from the
15 users of those computers, and by means of such conduct, obtained
16 the following approximate monies from the following advertising
17 service companies:

<u>COUNT</u>	<u>APPROXIMATE DATES</u>	<u>APPROXIMATE NUMBER OF PROTECTED COMPUTERS ACCESSED WITHOUT AUTHORIZATION</u>	<u>APPROXIMATE PAYMENT</u>
SEVEN	November 1, 2004 through November 19, 2004	26,975	\$4,044.26 from Gammacash
EIGHT	November 16, 2004 through December 7, 2004	8,744	\$1,306.52 from LOUDcash
NINE	January 15, 2005	19,934	\$2,888.11

Malice

- There are malicious people out there who want to:
 - Bring your system down and/or steal confidential data
- When attacker is a nation-state, attacks are far harder to stop
 - Many defensive techniques involve stopping attacks that have been seen before
 - But nation-states can use *new* attack vectors

Speed of Light

- Question: how long does it take light to travel from Berkeley to New York?
- Answer:
 - Distance Berkeley → New York: 4,125 km (great circle)
 - Traveling 300,000 km/s: 13.75 msec

Networking Latencies

- Question: how long does it take an Internet “packet” to travel from Berkeley to New York?
- Answer:
 - For sure ≥ 13.75 msec
 - In practice this boils down to ≥ 40 msec

Implications for Networking

- Question: how many cycles does your PC execute before it can possibly **get a reply** to a message it sent to a New York web server?
 - Answer:
 - **Round trip** takes ≥ 80 msec
 - PC runs at (say) 3 GHz
 - $3,000,000,000 \text{ cycles/sec} * 0.08 \text{ sec} = 240,000,000 \text{ cycles}$
- = An Eon**
- Communication **feedback** is always *dated*
 - Communication fundamentally asynchronous

Even a Problem for LANs

- Question: what about between machines directly connected (via a *local area network* or LAN)?

- Answer:

```
% ping www.icir.org
```

```
PING www.icir.org (192.150.187.11): 56 data bytes
```

```
64 bytes from 192.150.187.11: icmp_seq=0 ttl=64 time=0.214 ms
```

```
64 bytes from 192.150.187.11: icmp_seq=1 ttl=64 time=0.226 ms
```

```
64 bytes from 192.150.187.11: icmp_seq=2 ttl=64 time=0.209 ms
```

```
64 bytes from 192.150.187.11: icmp_seq=3 ttl=64 time=0.212 ms
```

```
64 bytes from 192.150.187.11: icmp_seq=4 ttl=64 time=0.214 ms
```

- 200 μ sec = 600,000 cycles
 - Still a loooong time ...
 - ... and asynchronous

Summary

- The Internet is a large complicated system that must meet an unprecedented variety of challenges
 - Scale, dynamic range, diversity, ad hoc, failures, asynchrony, malice, and greed
- An amazing feat of engineering
 - Went against the conventional wisdom
 - Created a new networking paradigm
- In hindsight, some aspects of design are terrible
 - But enormity of genius far outweighs the oversights

Next Lecture

- Read Sections 1.1-1.3 of the textbook
- Answer questionnaire
- Make sure you are on Piazza, bspace, etc.
- Remember to participate!
- Brush up on your Python
 - LearnStreet.com created by ex-122 students
 - Many other online resources....